

How to take the Dual of a Linear Program

Sébastien Lahaie*

October 1, 2008

The purpose of these notes is to: (1) explain how to recognize a linear program, (2) show how to take the dual of a linear program, and (3) list the fundamental results relating a linear program and its dual. These notes do not provide any proofs and do not explain any of the deep geometric insight behind linear programming duality. The main goal is to explain in detail the mechanical procedure of taking the dual.

There are other ways to take the dual than the one described herein, but this is my favorite. I find that it is readily memorized once it has been practiced on a few linear programs. It is perhaps slower and more tedious than other methods you might find in textbooks, but it is easy to memorize, and the intermediate steps produce useful information as a by-product.

1. Formulation

A linear program is a formulation of an optimization problem: a minimization or maximization of an objective function over some domain. The objective function is linear, and the domain, or *feasible set*, is defined by linear constraints. Rather than give the generic form of a linear program, here is a specific linear program that we will work with throughout this primer. Once we've studied this example, it should be quite clear what a linear program looks like in general.

$$\max_{x_1 \geq 0, x_2 \leq 0, x_3} v_1 x_1 + v_2 x_2 + v_3 x_3 \quad (1)$$

$$\text{such that } a_1 x_1 + x_2 + x_3 \leq b_1 \quad (2)$$

$$x_1 + a_2 x_2 = b_2 \quad (3)$$

$$a_3 x_3 \geq b_3 \quad (4)$$

The variables here are x_1, x_2, x_3 . The remaining terms are constants (e.g., v_1, a_2, b_3). A linear program consists of an objective function (1), and a set of inequality and equality constraints (2–4). The objective function

$$f(x) = v_1 x_1 + v_2 x_2 + v_3 x_3$$

is a *linear* function. Formally, this means that for two vectors x^1 and x^2 and real-valued constants c_1 and c_2 we have $f(c_1 x^1 + c_2 x^2) = c_1 f(x^1) + c_2 f(x^2)$. You can check that this

*. Yahoo Research, New York, NY, 10018. lahaies@yahoo-inc.com

indeed holds for our objective. The objective can either be a *maximization* or *minimization*. In our case, we have a maximization problem.

The left-hand side of each constraint is also a linear function. The right-hand sides are all constants. (You may encounter some linear programs where the right-hand side for some constraint includes some variables, but the constraint can always be rearranged so that the right-hand side is constant.) In a linear program, we cannot have strict inequalities; so $x_1 + x_2 < 3$ would not be allowed, for example.

Constraints that specify whether each variable is non-negative, non-positive, or unrestricted are special and we usually list those under the max or min. This only holds for these kinds of constraints. A constraint of the form $x_1 \geq 2$, for example, is not special, whereas $x_1 \geq 0$ is. In our linear program, x_1 must be non-negative, x_2 must be non-positive, and x_3 is unrestricted (it helps to be explicit about this).

2. Primal and Dual

The linear program we start with is typically called the “primal”. To each linear program there is associated another linear program called its “dual”. Deriving the dual from the primal is a purely mechanical procedure. Let’s see how to derive the dual of the linear program of the previous section. There are seven steps. The first two steps put the primal in a “standard form”.

Step 1. *If necessary, rewrite the objective as a minimization.*

In our case the objective (1) is a maximization, so we rewrite it as

$$\min_{x_1 \geq 0, x_2 \leq 0, x_3} -v_1x_1 - v_2x_2 - v_3x_3$$

A solution that maximizes an objective also minimizes the negative of that objective, so this does not change the set of optimal solutions to the program.

Step 2. *Rewrite each inequality constraint as a “less than or equal”, and rearrange each constraint so that the right-hand side is 0.*

After this step our linear program now looks as follows.

$$\begin{aligned} \min_{x_1 \geq 0, x_2 \leq 0, x_3} \quad & -v_1x_1 - v_2x_2 - v_3x_3 \\ \text{s.t.} \quad & a_1x_1 + x_2 + x_3 - b_1 \leq 0 & (5) \\ & x_1 + a_2x_2 - b_2 = 0 & (6) \\ & -a_3x_3 + b_3 \leq 0 & (7) \end{aligned}$$

Note that the special constraints listed below the min do *not* change.

Step 3. *Define a non-negative dual variable for each inequality constraint, and an unrestricted dual variable for each equality constraint.*

To constraints (5) and (7) we associate variables $\lambda_1 \geq 0$ and $\lambda_3 \geq 0$ respectively. To constraint (6) we associate λ_2 , which is unrestricted.

Step 4. *For each constraint, eliminate the constraint and add the term (dual variable)*(left-hand side of constraint) to the objective. Maximize the result over the dual variables.*

Concretely, for the first constraint (5) we would remove it and *add* the following term to the objective.

$$\lambda_1 (a_1 x_1 + x_2 + x_3 - b_1)$$

If we do this for each constraint (except of course special constraints), and maximize the result over the dual variables, we get

$$\begin{aligned} \max_{\lambda_1 \geq 0, \lambda_2, \lambda_3 \geq 0} \quad \min_{x_1 \geq 0, x_2 \leq 0, x_3} \quad & -v_1 x_1 - v_2 x_2 - v_3 x_3 \\ & + \lambda_1 (a_1 x_1 + x_2 + x_3 - b_1) \end{aligned} \tag{8}$$

$$+ \lambda_2 (x_1 + a_2 x_2 - b_2) \tag{9}$$

$$+ \lambda_3 (-a_3 x_3 + b_3) \tag{10}$$

It helps to think of this as a two-player game, with an “outer player” and an “inner player”. The outer player goes first and chooses some values for $\lambda_1, \lambda_2, \lambda_3$ (respecting the special constraints). With these values fixed, the inner player then chooses some values for x_1, x_2, x_3 so as to minimize the objective. Conscious that the inner player will do this, the outer player will choose values for $\lambda_1, \lambda_2, \lambda_3$ so that the minimum value in the inner player’s minimization problem is as large as possible.

Step 5. *We now have an objective with several terms of the form (dual variable)*(expression with primal variables), plus remaining terms involving only primal variables. Rewrite the objective so that it consists of several terms of the form (primal variable)*(expression with dual variables), plus remaining terms involving only dual variables.*

If we do this to the objective of the previous step, we obtain

$$\begin{aligned} \max_{\lambda_1 \geq 0, \lambda_2, \lambda_3 \geq 0} \quad \min_{x_1 \geq 0, x_2 \leq 0, x_3} \quad & -b_1 \lambda_1 - b_2 \lambda_2 - b_3 \lambda_3 \\ & + x_1 (a_1 \lambda_1 + \lambda_2 - v_1) \end{aligned} \tag{11}$$

$$+ x_2 (\lambda_1 + a_2 \lambda_2 - v_2) \tag{12}$$

$$+ x_3 (\lambda_1 - a_3 \lambda_3 - v_3) \tag{13}$$

This step takes great care. If you get any sign wrong or forget any term, the resulting dual will of course be wrong, but also very misleading and confusing.

Step 6. *Remove each term of the form (primal variable)*(expression with dual variables) and replace with a constraint of the form:*

- *expression ≥ 0 , if the primal variable is non-negative.*
- *expression ≤ 0 , if the primal variable is non-positive.*
- *expression = 0, if the primal variable is unrestricted.*

This step is not hard to memorize because it is in fact intuitive. Let’s consider term (11):

$$x_1 (a_1 \lambda_1 + \lambda_2 - v_1).$$

Because $x_1 \geq 0$, we introduce constraint $a_1 \lambda_1 + \lambda_2 - v_1 \geq 0$. Why must this hold? If we had $a_1 \lambda_1 + \lambda_2 - v_1 < 0$, then note that the inner player could choose $x_1 \rightarrow +\infty$ (i.e., arbitrarily

large), and thus the objective value is $-\infty$. Since the outer player wants to maximize the value of the inner player's minimization problem, it should therefore choose the values of $\lambda_1, \lambda_2, \lambda_3$ so that $a_1\lambda_1 + \lambda_2 - v_1 \geq 0$.

The same reasoning applies to the other two terms. The outer player must choose its values so that $\lambda_1 + a_2\lambda_2 - v_2 \leq 0$ otherwise the inner player can make the term

$$x_2(\lambda_1 + a_2\lambda_2 - v_2)$$

tend to $-\infty$ by choosing $x_2 \rightarrow -\infty$. Finally, since x_3 is unrestricted, the only way to make sure the term

$$x_3(\lambda_1 - a_3\lambda_3 - v_3)$$

cannot approach $-\infty$ is to choose the dual values such that $\lambda_1 - a_3\lambda_3 - v_3 = 0$. After these changes we have a new linear program. Note that the primal variables no longer appear.

$$\begin{aligned} \max_{\lambda_1 \geq 0, \lambda_2, \lambda_3 \geq 0} \quad & -b_1\lambda_1 - b_2\lambda_2 - b_3\lambda_3 \\ & a_1\lambda_1 + \lambda_2 - v_1 \geq 0 \end{aligned} \tag{14}$$

$$\lambda_1 + a_2\lambda_2 - v_2 \leq 0 \tag{15}$$

$$\lambda_1 - a_3\lambda_3 - v_3 = 0 \tag{16}$$

Step 7. *If the linear program in step 1 was rewritten as a minimization, rewrite the result of the previous step as a minimization; otherwise, do nothing.*

The result looks as follows. Optionally, the constraints can be also be rearranged in whichever form is most natural.

$$\begin{aligned} \min_{\lambda_1 \geq 0, \lambda_2, \lambda_3 \geq 0} \quad & b_1\lambda_1 + b_2\lambda_2 + b_3\lambda_3 \\ & a_1\lambda_1 + \lambda_2 \geq v_1 \end{aligned} \tag{17}$$

$$\lambda_1 + a_2\lambda_2 \leq v_2 \tag{18}$$

$$\lambda_1 - a_3\lambda_3 = v_3 \tag{19}$$

This completes the process of taking the dual. As an exercise, you can check for yourself that by taking the dual of this new linear program, you recover the primal.

3. Key Results

A linear program may be *infeasible*, *unbounded*, or have a *finite optimum*. A program is infeasible if there no solution that satisfies all the given constraints. For instance, suppose that in our original primal program we have $a_1 = a_2 = a_3 = 1$, and $b_1 = -1$ whereas $b_2 + b_3 \geq 1$. Then clearly there is no solution to constraints (2–4). (Feasibility only has to do with the constraints, not the objective.)

A linear program can alternatively be unbounded. This means that for any feasible solution, there is another feasible solution with strictly higher objective value, in the case of a maximization program, or strictly lower objective value, in the case of a minimization program. For instance, suppose that in our original primal we have $a_1 = a_2 = 1$ and

	Finite optimum	Unbounded	Infeasible
Finite optimum	Possible	Impossible	Impossible
Unbounded	Impossible	Impossible	Possible
Infeasible	Impossible	Possible	Possible

Table 1: Possible combinations for the primal and its dual.

$a_3 = -1$, whereas $b_1 = b_2 = b_3 = 0$, and finally the coefficients in the objective are all positive $:v_1, v_2, v_3 > 0$. Then note that $c(0, 0, 1)$ is a feasible solution for any $c \in \mathbf{R}$. The objective value of this solution is cv_3 , which we can make arbitrarily large by letting $c \rightarrow +\infty$. Technically, this means there is no optimal solution, even though there are feasible solutions.

If a program is feasible and bounded, it has a finite optimum (though it may not be unique). Table 1 lists the possible relationships between feasibility of a primal program and its dual. In particular, note that if the primal program is unbounded, the dual is infeasible, and if the dual is unbounded, the primal program is infeasible. But it may be the case that both are infeasible.

The value of the objective function of the primal program at an optimum is denoted V_P and the optimum value of the dual is denoted V_D . The major result of linear programming is the following.

Theorem 1 (Strong duality) *If a linear programming problem has an optimal solution, so does its dual, and $V_P = V_D$.*

This result relates the values of the primal and dual programs, but not their solutions. The following result is written in reference to our running example. It should be straightforward to adapt it to other linear programs.

Theorem 2 (Complementary slackness) *Let (x_1, x_2, x_3) and $(\lambda_1, \lambda_2, \lambda_3)$ be feasible solutions to the primal and dual problem, respectively. They are optimal solutions for the two respective problems if and only if*

$$\begin{aligned}\lambda_1(a_1x_1 + x_2 + x_3 - b_1) &= 0 \\ \lambda_2(x_1 + a_2x_2 - b_2) &= 0 \\ \lambda_3(-a_3x_3 + b_3) &= 0\end{aligned}$$

and

$$\begin{aligned}x_1(a_1\lambda_1 + \lambda_2 - v_1) &= 0 \\ x_2(\lambda_1 + a_2\lambda_2 - v_2) &= 0 \\ x_3(\lambda_1 - a_3\lambda_3 - v_3) &= 0\end{aligned}$$

These constraints are known as the *complementary slackness conditions*. Note that the first three constraints can be read off from (8–10), and the final three from (11–13). So we get the complementary slackness conditions as a by-product of our procedure for taking the dual.

The complementary slackness conditions are sometimes written in alternative forms. For instance, we can write the condition

$$\lambda_1 (a_1 x_1 + x_2 + x_3 - b_1) = 0$$

instead as

$$\lambda_1 > 0 \Rightarrow a_1 x_1 + x_2 + x_3 - b_1 = 0.$$

or equivalently (taking the contrapositive)

$$a_1 x_1 + x_2 + x_3 - b_1 < 0 \Rightarrow \lambda_1 = 0.$$

Given an optimal primal solution, the complementary slackness conditions define a system of equalities that can be solved to identify an optimal dual solution (the latter solution also has to satisfy the dual feasibility constraints), and vice-versa.